Please note that this is a pre-print version. For citation, please refer to the published version only.
The final publication is available at Springer via http://dx.doi.org/10.1007/s11229-008-9365-z.

A Plea for Pragmatics

Abstract Let *intentionalism* be the view that what proposition is expressed in context by a sentence containing indexicals depends on the speaker's intentions. It has recently been argued that intentionalism makes communicative success mysterious and that there are counterexamples to the intentionalist view in the form of cases of mismatch between the intended interpretation and the intuitively correct interpretation. In this paper, I argue that these objections can be met, once we acknowledge that we may distinguish *what determines the correct interpretation* from *the evidence that is available to the audience*, as well as from *the standards by which we judge whether or not a given interpretation is reasonable*. With these distinctions in place, we see that intentionalism does not render communicative success mysterious, and that cases of mismatch between the intended interpretation and the intuitively correct one can easily be accommodated. The distinction is also useful in treating the Humpty Dumpty problem for intentionalism, since it turns out that this can be treated as an extreme special case of mismatch.

0. Preamble

In contemporary philosophy of language, it is a commonplace that natural languages exhibit context dependence. Even those who argue that most of our expressions are semantically invariant across contexts agree that there is a certain class of expressions, namely *indexicals* (such as 'I', 'now', and 'here'), whose content is relative to the context of utterance. In order to accommodate this phenomenon, we may take the propositions expressed by sentences containing indexicals to be determined by sentence-context pairs, where the context may be represented by values of different parameters (such as agent, time, place, and so on). The question that now arises is what determines the values of the contextual parameters. On the

kind of view I will defend in this paper, the speaker's intentions determine these values, and thus the proposition that a sentence containing indexicals expresses in context depends on what the speaker intends in speaking. Let us call this view *intentionalism*.

My aim in this paper is neither to give a full intentionalist account, nor to provide positive reasons for accepting such an account.[1] Rather, I will restrict myself to a couple of objections against intentionalism that have been raised recently in the literature, and show how these can be met.[2] This means that a number of important issues, such as the specification of the content of the relevant intentions, fall outside the scope of this paper. However, my replies to the objections under consideration are independent of these issues.[3]

I will start by considering a general argument against intentionalism given by Christopher Gauker (2007), and argue that this argument fails. I will then go on to address a different, but related class of problems, which concern cases of *mismatch*, i.e., cases in which the intuitive referent differs from the one that the speaker's intention would determine. Such cases have been discussed frequently in the literature, and my suggestion is that there is a general strategy available for the intentionalist in treating these problems, which is closely related to the concerns raised in the discussion of Gauker's argument. Finally, I claim that the Humpty Dumpty problem for intentionalism can be viewed as an extreme case of mismatch, and thus that it can be treated along the same lines as other such cases.

1. Gauker's argument

In his paper 'Zero tolerance for pragmatics' Christopher Gauker argues for the following thesis: "*The proposition that a speaker's words express in context never depends at all on what the speaker intends in speaking.*" (Gauker 2007: xx. *Synthese Online First*, no page numbers.) Gauker uses the term 'pragmatics' to cover only the studies of the ways in which the proposition expressed depends on the speaker's intention, and thus denies that there

anything for pragmatics (in this sense) to study. This denial is incompatible with intentionalism.

Gauker distinguishes between two different ideas: that the intentions are *included* in the context and that the intentions *determine* the context. I will follow Gauker in taking the latter to be the correct idea. Since the correct interpretation of indexicals is determined by the context, and the concern in this paper is the interpretation of indexicals, it will be assumed throughout that the context determines the correct interpretation of a sentence (as uttered in that context) as well as the proposition expressed. Hence, on the intentionalist view, the speaker's intentions determine the context, and thereby they also determine the correct interpretation and the proposition expressed.

Gauker's argument has to do with the possibility of giving a plausible account of communication once we accept the intentionalist story about how the proposition expressed gets determined. Here is how he puts it:

> Hearers have little access to what people have in mind apart from the interpretation of what they say. So interpretation would be a problem that hearers could not solve if they had to know what a speaker intended in order to identify the content of the context and thus to interpret the speaker's utterance. So a hearer cannot reliably employ a method of interpretation that requires the hearer to have an independent insight into the speaker's intention. But, as we have seen, the proposition expressed by an utterance has to be one that a competent hearer could assign to it using a method of interpretation that the hearer could reliably employ on the basis of features of the situation that the hearer could normally be aware of. So the determinants of the context have to be features of the situation that are typically accessible to the hearer, and it cannot be the case that there are elements of the context that are determined by the speaker's intentions. (Gauker 2007: xx.)

The idea is clear enough. Since the speaker's intentions are not independently accessible to the audience, they cannot be what determine the correct interpretation. If they were, interpretation would not be able to proceed via a reliable method, and then it would be hard to see how the high amount of success ordinarily displayed in everyday communication could be explained. So, an intentionalist framework cannot provide a plausible account of communication, and should thus be rejected. The argument can be restated as follows:

> (1) The interpreter does not have knowledge of the speaker's intention independently of what proposition the speaker expresses.
> (2) An interpreter must be able to proceed via a reliable method.
> (3) In order for interpretation to proceed via a reliable method, the interpreter must have access to the determinants of the context, independently of what proposition the speaker expresses.
> (4) Hence, the determinants of the context cannot be the speaker's intentions.

It may seem like a neat and compelling argument. It is obviously valid, and the premises seem prima facie plausible. However, I think the argument fails. Gauker's mistake lies in assuming that the things that guide the hearer to the correct interpretation must also be the things that determine the correct interpretation. There is no need for intentionalists to accept this assumption; in fact, I think that what Gauker's argument shows is that they are more or less committed to a rejection of it.

2. Rebutting Gauker's argument

The intentionalist can accept (2); there is no reason to deny that interpretation should proceed via a reliable method, and intentionalism does not commit us to such a denial. The

intentionalist can also accept (1), since there is no need to claim that interpreters have independent access to intentions. Of course, the interpreter must have access to clues about the correct interpretation independently of the correct interpretation itself. But that does not mean that the interpreter must have independent access to the *determinants* of the context. There are other clues that may guide the interpreter to the correct interpretation. These include the words uttered, knowledge of a shared language, shared environment, shared conventions etc. The intentionalist should distinguish between what *determines* the context (and hence the proposition expressed) and what *guides the interpreter* to the correct interpretation. Of course, speakers will normally try to bring these into harmony, but that does not mean that they necessarily coincide. Once we acknowledge that this distinction could be drawn, we open up for the possibility that an interpreter may arrive at the intended interpretation via a reliable method without any prior grasp of the intention. Accordingly, I suggest that the intentionalist should reject (3).

I will not go on to *draw* this distinction, in the sense of giving a detailed account of the things falling on each side of it. It is enough for present purposes to observe that this distinction makes perfect sense from an intentionalist point of view. On the one side, we have the speaker intentions, which determine the context (and hence the proposition expressed). On the other side, we have the external evidence that is available to the audience. I grant that much more needs to be said about this in the course of constructing a more substantial intentionalist theory, but this falls outside the scope of this paper. Anyway, it should be clear that there is no reason to suppose that the intentionalist is committed to denying that there is such a distinction. Of course, if we assume that the external evidence determines the context, this distinction may not make sense, but that is irrelevant here. The objections I consider in this paper are *internal* in the sense that they aim to show that intentionalism has certain implausible consequences. This means that as long as the distinction makes sense from an

intentionalist point of view, we should be allowed to use it in replying to these objections. (Cf. note 3.)

So, how does this distinction help in supporting the denial of (3)? What we need is some reason to think that there could be a reliable method of interpretation that did not rely upon independent access to the determinants of the context. The suggested distinction allows us to give such a reason, and the basic idea may be summed up in a slogan: *There are two routes to the correct interpretation*. The context, remember, is represented by values of different parameters. One obvious way to reliably arrive at the correct interpretation would be to have direct access to the determinants of these values. That is the *first* route to the correct interpretation. Now, how could an interpreter reliably arrive at the values representing the context pertaining to the utterance, without having access to the determinants of these values? Well, she could take the external evidence into account, and infer the values of the parameters from this evidence. This is the *second* route to the correct interpretation. Given that the available evidence tends to guide the interpreter to the correct interpretation, the method will be reliable, although not foolproof.

An analogy may be helpful here. Consider a detective who sets out to solve a murder case. His task is to find out the values of three parameters: *who*, *where*, and *with what*. Suppose that the actual facts determine that the butler did it in the kitchen with the frying pan. These facts are what determine the correct solution, but since they occurred in the past, they are not directly available to the detective (since he did not witness the murder). Instead, he takes the available evidence into account. He is a very good detective, with an impressive track record, so we may assume that given that there is enough evidence of the right kind, he will reliably arrive at the correct values of the three parameters. His method is reliable, but it does not require independent access to the facts that determine the correct solution. The inference he makes proceeds via the external evidence, like stains of the victim's blood in the

kitchen, and the butler's fingerprints on the handle of the frying pan. It does not proceed via the determinants of the values. There is nothing mysterious about this kind of inference, and it seems plausible to take the detective's method to be reliable.

Now, consider the case of communication with indexicals. The interpreter's task is to find the value of the relevant contextual parameters. On the assumption that the speaker's intentions determine the context, the determinants of the values are not directly available to the interpreter. So, the first route to the correct interpretation is closed. Instead, she takes the second route, i.e., she takes the available external evidence into account, like the words and gestures used by the speaker, the shared environment, and so on. Supposing that she is a good interpreter, like most language users, and that there is enough of the right kind of evidence available, she is likely to infer the correct contextual values, and thus arrive at the correct interpretation. Of course, it is up to the speaker to provide adequate evidence for what the values of the contextual parameters are, and it is up to the audience to track the intended context by means of the evidence provided. According to this picture, the inference the interpreter makes proceeds via the external evidence, not via the determinants of the context. Just like in the case with the detective above, there is nothing mysterious about the inference, and there is no reason why we should not take the interpreter's method as we have described it to be reliable.

Of course, the connection between the determinants and the evidence is different in the two cases, but the important thing is that there is a reliable connection between them. Just as there is a close connection between the events involved in the murder and the evidence, there is, on the intentionalist account, a close connection between the speaker's intentions and the evidence she provides the audience with. Consider the following passage by Kent Bach:

> You do not just have something "in mind" and hope that your audience is a good mind reader. Rather, you decide to refer to something and try to select an expression whose utterance will enable your audience, under the circumstances, to identify what you are referring to. These circumstances are comprised of mutually believed matters of fact, such as what is in plain view to both of you, including any gestures on your part, as well as shared background information. (Bach 1992: 299.)

The idea is that, on the intentionalist view, the speaker's intention plays a role not only in that it determines the context, but also in that it affects the speaker's choice of communicative devices, and thus that there is a close connection between the intention and the evidence that the speaker makes available to the audience. This connection gives us reasons to think that proceeding via the evidence could be a reliable method for arriving at the interpretation determined by the speaker's intention.

On the picture I have sketched above, a reliable method of interpretation does not require independent grasp of the determinants of the context. As long as there is a sufficiently high degree of match between the evidence provided and the intended context, communication with indexicals is likely to succeed, even if we assume that intentions determine the context. What a normal speaker in normal circumstances will do is to make use of external factors in order to provide the audience with as much and as clear evidence as she believes is required for her utterance to be interpreted as intended.

Of course, nothing about this picture *shows* that the speaker's intentions determine the context (and thereby the proposition expressed). Moreover, I have left out many important aspects that a full intentionalist theory of indexicals must deal with, such as specifying the content of the relevant intentions. However, what I set out to do above was neither to construct a full intentionalist account, nor to give a strong positive case for intentionalism, but rather to question (3) by showing that there could be a reliable method of

interpretation that does not include independent access to the determinants of the context. The central idea is that there could be a route to the correct interpretation apart from the one that proceed over the determinants of the context. One can acknowledge that there could be such a route while leaving it entirely open whether or not intentions are the determinants of the context.

3. Cases of mismatch

Gauker illustrates his conclusion with a case of *mismatch* between what any competent speaker would suppose the expression used referred to and the intended reference. The story goes as follows. Harry is trying on ties, and Sally is giving advice. When Harry is wrapped in a garish pink-and-green tie, Sally remarks: 'That would go nicely with your new jacket.' In fact, Sally is thinking of the tie Harry tried on two ties back, and is intending to refer to that one rather than the one presently wrapped around Harry's neck. However, the only reasonable interpretation from Harry's point of view is that 'that' as used by Sally on this occasion refers to the pink-and-green tie he is presently wearing. Here is what Gauker says about this case:

> So if we said that the referent of Sally's demonstrative was the tie she intended to refer to, we could not maintain that the proposition her utterance expressed was a proposition that Harry could assign to it using a method of interpretation that he could reliably employ on the basis of features of the situation that he could normally be aware of. Instead, we should say that the reference of her demonstrative was the pink-and-green tie around Harry's neck. (Gauker 2007: xx.)

There are other similar cases in the literature, for instance concerning communication with post-it notes. Here is how it might go. John has called in ill, and in order to inform visitors that John will not be in his office that day, Bill puts a post-it note with the sentence 'I am not

here today' on what he mistakenly believes to be John's office. In fact, Janet has been the occupant of that office for over a month, and Bill is the only one who does not know this. Now, intuitively, according to the correct interpretation of the note, the 'I' on the note refers to Janet, despite the fact that Bill intended it to infer to John. Other cases, like the ones given by Marga Reimer (1991), concern uses of demonstratives where there is some reason to think that the speaker's intention is overridden by external factors. For instance, in the case where I say 'These are my keys' while mistakenly grabbing my officemate's keys, the intuitive reference of 'these' is the keys I am *actually grabbing*, not the keys I *intend* to refer to, i.e., my own keys. (Reimer 1991: 190-191.)

Cases like these are supposed to make us suspicious about the idea that the speaker's intentions determine the context, by showing that it leads to counterintuitive results; there is a mismatch between the intuitive referent (if there is one) and the one that would be determined by the relevant intention. But if intentions do not determine the context, and thereby the referent, what does?

According to *conventionalism*, it is conventions rather than intentions that determine the reference of indexicals. For instance, Corazza et al. (2001) argues that in order to account for cases like the one with the post-it note, we should take the value of the agent parameter to be "given by the social or conventional setting in which the utterance takes place." (Corazza et al. 2001: 13.) So, on the conventionalist view, there is a convention to the effect that 'I' on a post-it note attached to someone's office door refers to the usual occupant of that office. This yields the intuitively right reference, in contrast to intentionalism, which seems to end up with a mismatch. (Cf. Corazza et al. 2001: 13.)

An apparent problem with the conventionalist picture, raised by Stefano Predelli, is that it seems to lead to a proliferation of conventions. For instance, we seem to need different conventions for 'today' as used on post-it notes saying 'I am not here today'

and 'Today the dean is getting on my nerves', where the former is written on Monday to inform readers of the note on Tuesday that I am not in (on Tuesday), and the latter is written on Monday only to record my annoyance (on Monday). (Predelli 2002: 314.) Moreover, Predelli suggests that even if there were different conventions in play, it would be the intentions of the writer of the note that determined which convention she appealed to. However, as pointed out by Jonathan Gorvett (2005: 307), it is not clear that Predelli's understanding of these notes is the correct one. Given that we do not presuppose that the writer's intentions are relevant, both could plausibly be interpreted as referring to the day when the note is read. Thus it is unclear what this example really shows. Moreover, Gorvett argues that conventions are established regularly among language users, so the proliferation of conventions is not really as implausible as Predelli thinks. (Gorvett 2005: 306-307.)

Another way to make reference depend on external factors is to follow Gauker (2007) in claiming that *accessibility criteria*, such as salience, prior reference, relevance, and charity determine the context pertaining to the utterance, which in turn determines the reference of the indexicals and demonstratives used in the utterance. Since a proper evaluation of these views falls outside of the scope of this paper, I will not go into any deeper discussion regarding these alternatives. What is certainly correct about them is that external factors play an important role in communication. But this is not something that the intentionalist must deny, even though she must deny that the external factors are the determinants of the context (and thereby the proposition expressed).

Firstly, it is perfectly consistent with the intentionalist picture that these external factors are used as *devices* for communication. As I stressed above, the external factors are important in the account of communication, even on the intentionalist account, since the speaker needs to *exploit* these in order to get her message across. Secondly, I suggest that the intentionalist should take the external factors to determine *standards* for judging how

reasonable or legitimate a given interpretation is. I will refer to these as the *external standards*. Normal cases of communicative success will presumably be cases where the there is little (or no) mismatch between these external standards and the intention; normal cases of communicative failure will presumably be cases where there is some kind of mismatch. Intentionalism is compatible with the existence of both kinds of cases.

So, how should the intentionalist respond to the challenge raised by cases of mismatch? One strategy, adopted by Bach in his reply to Reimer, is to argue that the appearance of mismatch is due to a failure to appreciate the nature of the referential intention. Reimer takes the case with the keys to be a counterexample to intentionalism, but Bach protests that the relevant intention is not the intention to *refer to my keys*, but rather the intention to *refer to the keys I grab*, since the audience "was to identify what I was using the word 'these' to refer to by thinking of the keys, not as my keys but as the keys I grabbed." (Bach 1992: 296.) Bach's reply is thus that in this case there is no genuine mismatch, and that we can come to realise this if we spell out the intention properly.

I will suggest a different, more general strategy. However, that does not mean that I object to Bach's treatment of this case. A proper specification of the relevant kind of intention in cases like these is an important part of a full intentionalist account. Indeed, such a specification would put us in a position to judge whether or not there is a *genuine* mismatch in the case at hand, and if there is not, then it does not pose any problem for intentionalism. But suppose that we were to come across genuine cases of mismatch between the intended reference and the intuitive referent, or between what the speaker intends to refer to and what the audience could reasonably or legitimately take the reference to be? Would that refute intentionalism?

According to my suggestion, the answer is no. Intentionalists can acknowledge the possibility of genuine cases of mismatch as long as they are careful to distinguish between

the *correct interpretation* (or the proposition expressed) and the *interpretation the audience may legitimately take to be the correct one* (or the *reasonable interpretation* for short). Intentionalism entails that the former is determined by the speaker's intentions, but there is nothing about this view that prevents the latter from being determined by the *external standards*. The idea is that our intuitions about genuine cases of mismatch track the *reasonable* interpretation rather than the *correct* interpretation.

Thus, the possibility of genuine cases of mismatch between what the audience could reasonably take the reference to be and what the speaker intends it to be is not a problem for the intentionalist. On an intentionalist view, these may well come apart, as long as the correct interpretation is distinguished from the reasonable one. Given this distinction, such a mismatch does not entail a mismatch between the intended interpretation and the correct one. In cases where there is a genuine mismatch between the speaker's intentions (which determine the correct interpretation) and the external standards (which determine the reasonable interpretation), communication is likely to fail, but that is just what we should expect. This applies straightforwardly to Gauker's tie case. On this occasion, Harry's method of interpretation yields an incorrect interpretation, since it leads him to an interpretation different from the intended one. Of course, that does not show that his method is not reliable in general, it only shows that it is not foolproof. Again, in cases like these, when the speaker provides misleading evidence, we should expect communication to fail. Even if the external evidence does not *determine* the correct interpretation, it is all that the interpreter has to go on.

If a speaker is not well enough informed about the external factors which she relies on in providing the audience with evidence for the intended interpretation, she will typically use the external devices in a defective way. As a result, the intended interpretation and the reasonable interpretation come apart. In such cases of mismatch, communication is likely to fail, and the speaker is to blame. But this does not show that the correct interpretation

is the one we (not the speaker) would expect the audience to arrive at, rather than the one intended by the speaker. The upshot is that the intentionalist can claim that our intuitions about these cases track something different from the correct interpretation, namely the interpretation the audience may legitimately take to be the correct one. A possible explanation for this is that we might tend to sympathise with the interpreter rather than the speaker in cases where the speaker is to blame for the communication failure.

Of course, merely distinguishing between the correct interpretation and the interpretation the audience may legitimately take to be the correct one is not to give any positive reason to think that speaker intentions determine the context (and hence the correct interpretation of indexicals). Such a distinction should be recognised by everyone anyway, in order to allow that reliable methods sometimes yield the wrong results. My point is rather that by taking this distinction into account, we can explain away intuitions to the effect that the correct interpretation sometimes deviates from the intended one by claiming that those intuitions track the reasonable rather than the correct interpretation. Or to put it in terms of contexts, the existence of cases in which the context determined by the speaker's intentions seems to differ from the one which we could reasonably expect to pertain to the utterance, does not show that the context pertaining to the utterance is not determined by the speaker's intentions.

The main idea may be summed up as follows. Our judgements about which interpretation is the reasonable one depend on the evidence provided by the speaker. Once we distinguish the evidence provided from the determinants of the context, we can also distinguish the correct interpretation from the reasonable one. In light of these distinctions, cases of mismatch between intended interpretation and reasonable interpretation are quite compatible with the claim that the correct interpretation is the intended one. Cases in which the proposition expressed intuitively differs from the one intended by the speaker can be

explained away by pointing out that our intuitions in those cases track the reasonable interpretation rather than the correct interpretation.

4. The Humpty Dumpty problem

Let us now turn to what may be called the Humpty Dumpty problem for intentionalism. The problem, recently discussed by Corazza et al. (2002), Predelli (2002), and Gorvett (2005), is a familiar one: Intuitively, speakers cannot use indexicals to refer to anything they like. For instance, it seems that John cannot, while standing on Trafalgar Square at 2 p.m. on the 15th of June 2007, use the sentence 'I am here now' to express the proposition that Bill was in Singapore at 5 p.m. on the 23rd of January 2006. But as long as there are no constraints on reference apart from the speaker's intention, and no constraint on intention forming, intentionalism entails that John *can* express this proposition by using those words, since the intention is what determines the reference. This looks like bad news for intentionalism.

There are similarities, but also important differences, between the response I will suggest and the one given by Predelli (2002). In order to illustrate this, I will briefly discuss the exchange between Gorvett (2005) and Predelli, and give my view on the matter. Predelli accepts, for the sake of argument, that one can intend an utterance to be interpreted with respect to any context, but denies that it follows from this that one can *use* indexicals to refer to pretty much anything. He distinguishes between what the indexicals *refer* to in the sense of what they are semantically associated with in virtue of the mechanisms governing expressions of this kind, and what the indexicals can be *used to refer to*, in the sense of what information a speaker can convey by employing them. (In the following paragraphs, I will sometimes refer to these senses as the *first notion of reference* and the *second notion of reference*, respectively.) His reply is then that the intuition that lies behind the Humpty Dumpty problem is to be understood as the claim that indexicals cannot be *used to refer* to

anything whatsoever, i.e., as the claim that "one does not always manage to convey contents pertaining to whatever individual or time she wishes." Predelli finds this claim plausible, but points out that it is perfectly consistent with intentionalism. (Predelli 2002: 314-315.)

Gorvett (2005: 300) complains that Predelli's first notion of reference has no obvious connection to communication, and thus is uninteresting in the present context. Hence, we should go for the second notion of reference. He goes on to argue that the second notion is not fit for an intentionalist account, since the most plausible account of what determines what the audience *actually* grasps is one given in terms of conventions, and this "removes the power of the intention from the picture." (Gorvett 2005: 305.) I think there is something right about Gorvett's remarks, but I do not think that he has succeeded in refuting intentionalism. Let me explain.

Firstly, I think we should stick with the first notion of reference. What is right about Gorvett's remark is that this notion does not carry an obvious connection to communication on its sleeves, and Predelli does not say anything about how to fit it into a theory of communication. However, that does not mean that this cannot be done. The problem with Gorvett's criticism is that he assumes that any notion of reference that can play a role in an account of communication must have an implausibly strong connection to communicative success. Here is what he finds problematic about Predelli's first notion of reference:

> For Predelli reference is simply a relationship between a word and an object but one that does not require or entail successful communication. That a word refers to an object does not mean that it can necessarily be used to communicate a thought about that object. (Gorvett 2005: 300)

Certainly, we need to fit our preferred notion of reference into our theory of communication, but why should we demand that reference "require or entail successful communication"? On

the intentional picture sketched above, a normal speaker who wants to say something about a certain object will decide to refer to it, and then choose whatever devices she believes will help her getting her message across. Here the external factors become relevant, and if she fails to make the right kind of evidence available to the audience, communication is likely to fail. In these situations, she will not *use* the expressions to refer in Predelli's sense, since this would entail that the audience grasped her message. However, why should we deny the possibility of situations in which the speaker succeeds in referring, but fails in getting her message across? This is just what we should expect in genuine cases of mismatch.

Secondly, I agree with Gorvett that the intentionalist should say something about the limitations on the forming of the relevant kind of intentions. Gorvett suggests that the intentionalist should argue that it is in some appropriate sense impossible for speakers to form certain intentions. Predelli does not take this route, but chooses to bypass the issue of what limitations there are on speakers' intention forming by accepting, for the sake of argument, that any context can be intended, and thus determined as the one pertaining to the utterance. Again, I think that something should be said about these limitations. However, I do not think that we need a *general* claim about what kind of intentions it is *possible* or *impossible* for speakers to form – a weaker claim about limitations on speakers' abilities to form referential intentions would be quite enough.[4] My suggestion is that it is in general *hard* for speakers to refer in the non-standard ways described in the Humpty Dumpty problem, since they are *used to* conform to the external standards (at least to some extent), and in fact *normally* refer (roughly) in accordance with it.

It may not be *impossible* to refer in this deviant way. However, speakers normally have a very strong inclination to conform to the external standards. It may even *seem* impossible for normal speakers to deviate from it in this extreme way. On the one hand, if a speaker *believes* that it is impossible to refer in a certain way by using certain indexicals,

she will not be able to form the intention to refer in that way. In that sense, *some* referential intentions are *sometimes* impossible for *some* speakers to form. On the other hand, if a speaker believes that it is possible to refer in a certain deviant way, there is no reason to deny that she could form an intention to do so. In such cases, there will be a mismatch between the intended interpretation and the reasonable interpretation. In extreme cases, like the one just described, it will be virtually impossible to arrive at the correct interpretation, unless we are familiar with John's deviant linguistic habits. The success rate of our methods of interpretation depends on the degree of match between the correct interpretation and the external standards, so it should not surprise us that communication will break down in these extreme cases, where there is a high degree of mismatch. But, again, this does not mean that a view according to which such cases of mismatch are possible renders our ordinary methods of interpretation unreliable.

In effect, then, I suggest that the intuition behind the Humpty Dumpty problem is basically of the same kind as in other cases of apparent mismatch between intended interpretation and correct interpretation. If this is right, we should be able to make use of the distinction between correct interpretation (which is determined by the intention) and reasonable interpretation (which is determined by the external standard) in the treatment of the Humpty Dumpty problem as well. Here is how it might go: In cases like the one just described, the deviant speaker, in this case John, intends to refer in a way that is completely different from what the external standards require. We reasonably take John's utterance to refer in accordance with the external standard, and thus take the proposition expressed to be that John is on Trafalgar Square at 2 p.m. on the 15th of June 2007. However, John refers in a highly idiosyncratic way and thus in fact expresses the proposition that Bill was in Singapore at 5 p.m. on the 23rd of January 2006. We would expect communication to fail in cases like this, and we would also expect normal speakers to find it hard to deviate from the external

standards in this extreme way, but there is no need to conclude that John *could* not refer in his idiosyncratic manner. Our intuitions to that effect could be explained by the fact that *normal* speakers would not be able to do this, since they are mentally constrained by their linguistic habits.

The defence presented here is similar to Predelli's in that we both attempt to explain away the intuition behind the Humpty Dumpty problem by showing that its plausibility could be explained in a way that is compatible with intentionalism. However, as I hope the above discussion illustrates, there are some important differences. In particular, Predelli accepts, for the sake of argument, that indexicals can refer in any way the speaker likes. My reply, on the other hand, indicates how the speaker's beliefs about external standards can help explain why she cannot intend to refer to anything whatsoever.

5. Conclusion

To sum up, then, we have seen that intentionalism can be defended by appeal to a distinction between the determinants of the proposition expressed, and the external standards by which we judge whether or not a given interpretation is reasonable. Even if the determinants of the context (and the proposition expressed) are not directly available to the interpreter, she can still arrive at the correct interpretation via a reliable method, as long as there is a sufficient degree of match between the correct interpretation and the external evidence provided. This move undermines Gauker's argument for his zero tolerance view, and it also shows how the intentionalist can explain away certain allegedly problematic cases involving mismatches between the intended interpretation and the reasonable one. Finally, we have seen that the intuition behind the Humpty Dumpty problem can be accounted for in terms of normal speakers' strong inclination to conform to the external standards. Normally, speakers try to avoid mismatches between their referential intentions and the external standards, and their

linguistic habits make it hard for them to use indexicals to refer in highly deviant ways. However, that does not mean that it is in general impossible for them to deviate, and if they do, our (normally reliable) methods of interpretation are likely to lead us astray. This should not strike us as problematic, as long as we do not require our methods of interpretation to be foolproof.

Please note that this is a pre-print version. For citation, please refer to the published version only.
The final publication is available at Springer via http://dx.doi.org/10.1007/s11229-008-9365-z.

NOTES

1. One positive reason for accepting intentionalism is that merely external factors, such as the spatio-temporal position of the speaker seem insufficient. They will not determine how close to the speaker something would need to be in order to fall under the predicate 'is here' or which time span is the extension of 'is here'. According to intentionalism, whether or not something falls in the extension of 'is here' depends on whether or not the position of this object is included in the area intended by the speaker.

2. We may distinguish between a stronger and a weaker intentionalist claim. The stronger claim is that intentions are the sole and ultimate determinants of the content of indexicals. The weaker claim is that intentions play a role in determining the content. The latter, but not the former is compatible with the idea that other (external) factors play a crucial role in determining the content. Some of the objections I consider, for instance the Humpty Dumpty objection I address at the end, are mainly directed at the stronger claim, while Gauker's objection is directed at the weaker claim. However, since I want to defend the stronger claim as well, and the stronger claim entails the weaker claim, I will not bother making this distinction in what follows.

3. The objections I consider are of the following form: If intentionalism were true, then *p* would hold, where *p* is some implausible proposition. There is no harm in supposing the antecedent to be true in the course of replying to these objections, since the aim is merely to undermine the conditionals.

4. I owe this insight to *****.

References

Bach, K. (1992): 'Paving the Road to Reference.' *Philosophical Studies* 67: 295-300.

Corazza, E., Fish, W., and Gorvett, J. (2002): 'Who is I?' *Philosophical Studies* 107: 1-21.

Gauker, C. (2007): 'Zero Tolerance for Pragmatics.' *Synthese* (Online First). (Published online 12th of June 2007, no page references.)

Gorvett, J. (2005): 'Back through the Looking Glass: On the Relationship between Intentions and Indexicals.' *Philosophical Studies* 124: 295-312.

Predelli, S. (1998a): 'Utterance, Interpretation, and the Logic of Indexicals', *Mind and Language* 13: 400–414.

Predelli, S. (1998b): 'I Am Not Here Now.' *Analysis* 58: 107–15.

Predelli, S. (2002): 'Intentions, Indexicals and Communication.' *Analysis* 62: 310–316.

Reimer, M. (1991): 'Demonstratives, Demonstrations, and Demonstrata.' *Philosophical Studies* 63: 187-202.